By Bill Travis, Senior Technical Editor

WITH HOT-SWAP AND LOAD-SHARE POWER-MANAGER ICs, YOU CAN FORGET ABOUT SUPPLIES AND CONCENTRATE ON IMPORTANT SYSTEM CONSIDERATIONS.

Hot-swap, load-share ICs protect circuits and supplies

IN SYSTEMS THAT USE multiple power supplies, hot swapping and current sharing are important considerations. Hot swapping refers to the insertion or removal of modules, boards, or power supplies into or from powered-up buses or connectors. PCMCIA, Universal Serial Bus (USB), CompactPCI, SCSI, and other standards provide strict mandates on inrush current, power-supply sequencing, and other power-up and -down aspects. Current, or load sharing involves equalizing, insofar as possible, the currents from multiple power supplies connected in parallel. Several recent ICs provide hot-swapping and current-sharing management functions, eliminating the need to design discrete circuitry for protection and power integrity in your system.

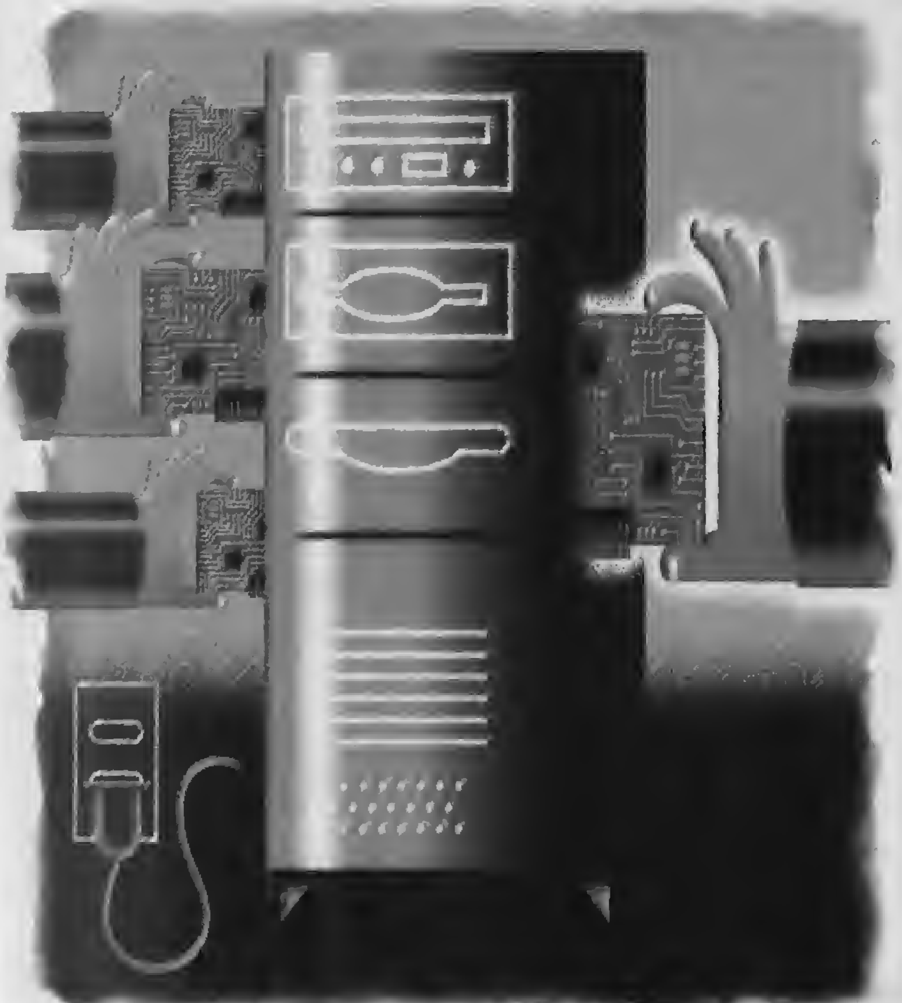

Hot swapping can be a scary proposition (**Reference 1**). For example, if you insert a circuit board into a live backplane, the board's large bypass capacitors can draw inrush currents as large as 100A from the backplane's power bus. Unrestricted, these currents can destroy the board's bypass capacitors, metal traces, or even its connector pins. The inrush currents also produce a droop in the supply voltage to the rest

AT A GLANCE

- ▹ Excessive inrush currents or improper supply sequencing can zap delicate circuitry.
- ▹ Designing your own hot-swap protection steals time and effort from the meaty design challenges in your system.
- ▹ Live-insertion/removal capability in standards such as Universal Serial Bus and CompactPCI is not a luxury; it's a necessity.
- ▹ Equitable current sharing in multiple-supply systems minimizes thermal, stability, and reliability compromises.

TABLE 1—UNITRODE HOT-SWAP POWER-MANAGER ICs

Devices with internal pass MOSFET			
Device	**Voltage range (V)**	**Maximum current (A)**	**$R_{DS(ON)}$ (Ω)**
UCC3912	3 to 8	4	0.15
UCC3915	7 to 15	5	0.15
UCC3916	SCSI, 2, 3	2	0.22
UCC3918	3 to 6	5	0.075
UCC3920	−3 to −15	4	0.1
Devices with external pass MOSFET			
Device	**Voltage range**		
UCC1914	4.5 to 35		
UCC3917	7 to >1000		
UCC3919	3 to 8		
UCC3913	−7 to >−1000		
UCC3921	−3 to >−1000		

of the circuit boards, which are connected to the backplane. The resulting out-of-regulation supply voltage can cause the affected boards to reset.

Most multiple-supply systems require the supply voltages to power up in a specific sequence. The sequence reverses upon power-down, with one supply kept alive long enough for software to control the power-down routine. Hot-swap power-management ICs provide both inrush protection and user-programmable supply sequencing. They also protect systems against power faults, and they monitor and signal the status of supply voltages and currents.

You can provide adequate hot-swap control, with perhaps a 90% success rate, by configuring controlled-turn-on FETs. RC-network delays at the gates of the FETs can produce the desired power-up delays. However, this approach has two potential problems: One is the long hours of tweaking that discrete analog circuits require. Although these hours may be enjoyable for some and relatively short for analog veterans, they are hours better spent on more mission-critical problems. The other problem is that 10% of the time, the solution doesn't work (**Reference 1**).

APPLICATION SPECIFICITY

Many hot-swap power-management ICs serve general-purpose applications; others are application-specific. Application-specific devices target specific buses or standards, such as USB, CompactPCI, or the various flavors of SCSI. They handle the voltage and current levels that the given standard mandates, and they provide norm-specified protection, timing, and sequencing features. Unitrode Corp offers a broad spectrum of hot-swap managers (**Table 1**).

Most of the hot-swap managers in **Table 1** are general-purpose devices; you choose your device according to the voltage and current requirements of the load. You can use the other devices in a variety of applications, but they are particularly useful in certain systems. The $2.95 (1000) UCC3915, for example, is intended for 12V hot-plugging applications, such as redundant array of independent disks (RAID) and popular communications ports, such as Ethernet and the advanced universal interface. Similarly, the $2.95 (1000) UCC3920 serves communications ports that use −12V rails. The UCC3916 targets SCSI Termpower systems and offers SCSI, SCSI-2, and SCSI-3 compliance.

Figure 1 shows the circuit architecture of the UCC3920. The other Unitrode hot-swap managers with internal pass MOSFETs have similar topologies with individual variations. The 3-bit DAC in the UCC3920 lets you program the fault-level trip current using a 3-bit digital code. The

Figure 1

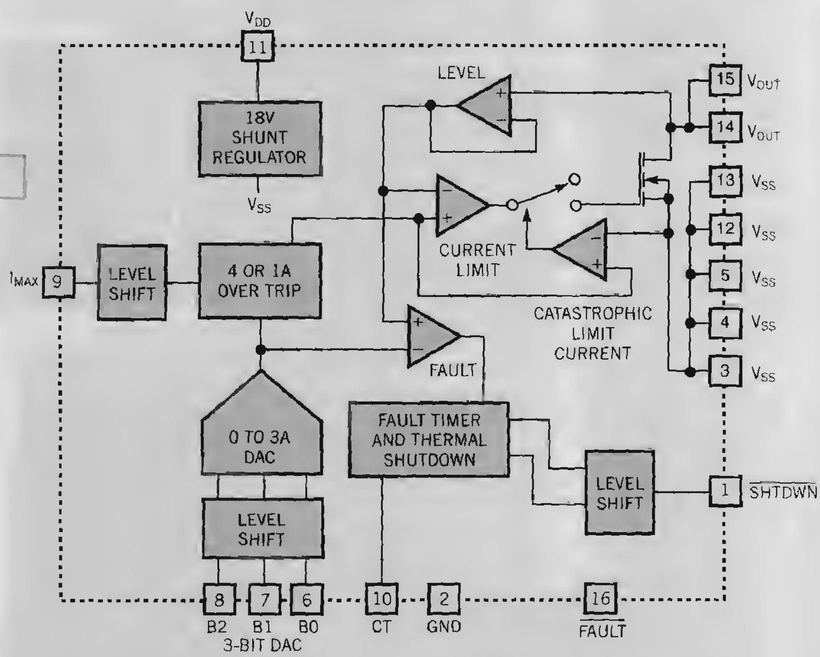

A 3-bit DAC in Unitrode's UCC3920 hot-swap manager allows you to program the fault-level trip current that suits your system.

programming resolution is 250 mA for currents of 0 to 500 mA, and 500 mA for currents of 500 mA to 3A. The UCC3920 handles −3 to −15V supply levels; its UCC3912 and 3915 siblings accommodate positive-voltage ranges of 3 to 8 and 7 to 15V, respectively. These managers incorporate a 4-bit DAC to allow fault-level programming of 0 to 3A with 250-mA resolution. The $3.05 (1000) UCC3918 uses an external resistor rather than a DAC to provide linear fault-level programming of 0 to 4A.

The devices in **Table 1** that use external pass MOSFETs include charge-pump circuitry, which generates the gate voltages needed to turn on the n-channel MOSFETs. The ICs would be simpler and use less silicon if the external device were a p-channel MOSFET. However, through a sad and immutable quirk of physics, a p-channel device uses about three times the silicon area of an n-channel MOSFET with the same ratings. It's more economical to connect a few capacitors and use an on-chip charge pump to drive an n-channel FET. **Table 1** gives no current ratings for the external-FET devices because that parameter is strictly a function of the MOSFET you specify.

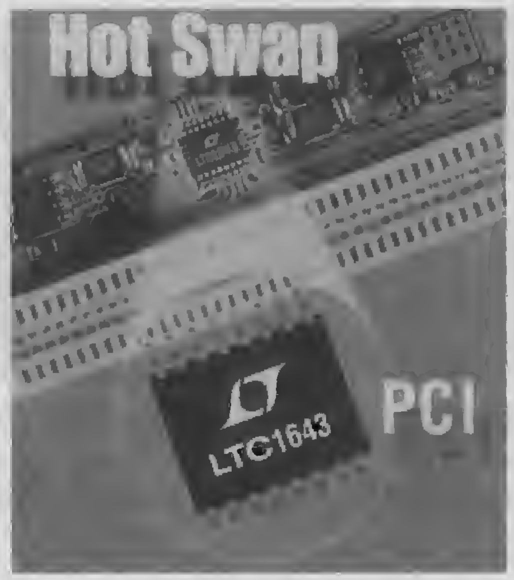

The LTC1643 hot-swap controller from Linear Technology makes live insertion/removal a safe operation in CompactPCI systems.

Figure 2 shows an elementary block diagram of a typical hot-swap IC that drives an external MOSFET. The charge-pump circuitry uses internal diodes and the external capacitors to configure a voltage tripler. V_{PUMP} provides plenty of head room to power the internal gate driver. I_{MAX}, the maximum sourcing current, is a programmable parameter. When the sensed output current is lower than fault level, the output MOSFET remains on. When the

CURRENT SHARING ENHANCES PERFORMANCE AND FAULT TOLERANCE

Robert Pauplis, Senior Application Engineer, Vicor Corp

Power architects often employ multiple power supplies or power converters to increase output power or provide fault tolerance. You can usually increase power by paralleling two or more converters. Fault-tolerant systems also use power supplies in parallel, but they may use ORing diodes to isolate modules during an output fault.

Current sharing is a necessity when paralleling supplies to increase power. Although 2N-redundant supplies generally require no current sharing, current sharing improves system performance. N+1-redundant supplies require current sharing. Current sharing and how best to achieve it are important issues in managing multiple power supplies or power converters.

Current sharing brings a number of benefits to parallel arrays. It improves transient response because each of the N converters shares the load step in a 1/N proportion. Each supply ages the same amount if equal current sharing exists, and thermal constraints are usually less demanding on systems using multiple supplies. Modular converters connected in parallel to increase power often find use in redundant applications, such as scalable power systems. The module array is readily expandable and may provide hot-swap capability.

You can implement current sharing in several ways; each method has advantages and disadvantages. Driver/booster arrays for the expansion of power usually contain one intelligent module or driver and one or more power-train-only modules or boosters. The driver sets and controls output voltage, and booster modules increase output power to meet system requirements. The advantages are that the array has only a single control loop, and it provides excellent transient response because no loop-within-a-loop stability issues exist. An advantage of driver-booster arrays is that load sharing is accurate, even during dynamic load conditions. Driver-booster arrays using only one driver do not support redundant operation.

The "droop-share" approach to current sharing employs a resistance in series with the load or an active circuit that allows the output voltage to drop in response to increasing load. The droop-share circuit has the advantages of simplicity and low cost. However, it is limited in application because it usually requires manual adjustment of the output voltage to achieve current sharing. Also, the series resistance degrades output-voltage regulation in droop-share circuits.

DC-coupled, single-wire paralleling involves two or more identical modules, each containing a circuit that monitors the current each supply delivers. This circuit actively adjusts the output voltage of each supply so that the multiple supplies deliver equal currents. However, this method has a number of disadvantages. Multiple control loops can cause stability problems, and they can provide poor transient response when a module fails. The method is also susceptible to single-point failures, which can defeat current sharing and, at worst, can cause a chain of failures.

A new fault-immune power architecture uses a digital current-sharing signal. Because it is an ac signal, dc blocking eliminates offsets in the system. Such modules allow designs to achieve high levels of availability and reliability. Additional advantages include excellent transient response, a high degree of immunity from system noise, and no loop-within-a-loop control problems.

current exceeds the fault level but is lower than the maximum sourcing current that I_{MAX} sets, the output remains on, and a fault timer begins to charge C_T. When C_T charges to a certain level, the output MOSFET turns off, and C_T slowly discharges. When C_T discharges to a certain low level, the hot-swap IC performs a retry, and the output MOSFET switches on again. The UCC3914 offers two distinct reset modes. In one mode, the IC repeatedly tries to reset itself if a fault occurs. In the other mode, when a fault occurs, the output stays off until you toggle a pin or until you turn the IC off then on again.

FLOATING ON AIR

Glancing at the 1000V-rated hot-swap ICs in **Table 1**, you might think that Unitrode uses an exotic process such as dielectric isolation to make its 1000V devices, but it does not. In a process called "bootstrapping" (derived from "pulling yourself up by your own bootstraps"), the $1.54 (1000) UCC3917 floats at the supply voltage and uses charge-pump circuitry to generate its supply voltage. The negative-supply, $1.40 (1000) UCC3913 and $1.42 (1000) UCC3921 use the negative input voltage as reference and derive their power from an external resistor connected to ground. The 1000V limit for these ICs is not etched in stone; the maximum voltage is a function of the ratings of the external components.

Designed for RAID systems, Harris Semiconductor's HIP1012 dual-voltage hot-swap controller safely applies and removes power to and from a disk drive, allowing you to swap drives during service operations. The IC provides active current limiting for your pin-selectable choice of dual-supply power buses: 5 and 12V or 5 and 3.3V. In a typical application diagram, the two current-sensing resistors, R_{SENSE}, provide inputs for the IC's overcurrent-protection circuitry (**Figure 3**). When the current through either resistor exceeds the user-programmed value, the controller enters its current-regulation mode, and the time-out capacitor, C_{TIM}, starts charging. When C_{TIM} charges to a 2V threshold, the IC switches off the n-channel MOSFETs.

The time-out period ranges, for example, from 4.4 msec with C_{TIM}=22 nF to 20 msec for C_{TIM}=100 nF. If a fault is three times the current-limit level, the controller turns off the MOSFETs in less than 3 µsec before entering its time-out period. The turn-off time is less than 1 µsec for a dead short. A rising edge on either power-on pin

Figure 2

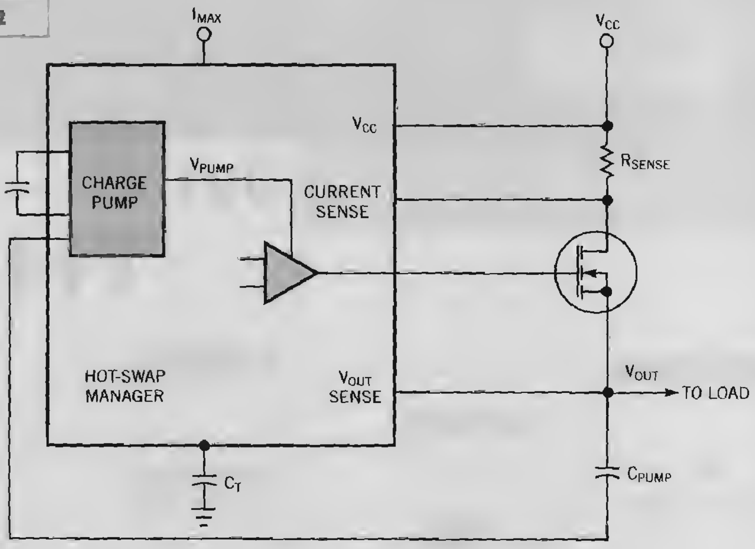

Bootstrapping via C_{PUMP}, an on-chip charge-pump circuit, provides the gate drive for the external MOSFET used with most hot-swap controllers.

FOR MORE INFORMATION...

For free information on power-management ICs such as those described in this article, circle the appropriate numbers on the postage-paid Information Retrieval Service card or use *EDN*'s InfoAccess service. When you contact any of the following manufacturers directly, please let them know you read about their products in *EDN*.

Harris Semiconductor
Melbourne, FL
1-407-729-5973
fax 1-407-729-5321
www.semi.harris.com
Circle No. 326

Linear Technology Corp
Milpitas, CA
1-408-432-1900 ext 2453
fax 1-408-434-6441
www.linear-tech.com
Circle No. 327

Texas Instruments Inc
Dallas, TX
1-800-477-8924 ext 4500
www.ti.com/sc/power
Circle No. 328

Unitrode Corp
Merrimack, NH
1-603-424-2410
fax 1-603-424-3460
www.unitrode.com
Circle No. 329

Vishay Siliconix
Santa Clara, CA
1-408-988-8000
fax 1-408-567-8950
www.siliconix.com
Circle No. 330

SUPER CIRCLE NUMBER
For more information on the products available from all of the vendors listed in this box, circle one number on the reader service card. *Circle No. 331*